

**IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
WACO DIVISION**

RUBRIK, INC.,)	
)	
Plaintiff,)	
)	Civil Action No. 6:21-cv-00135-ADA
v.)	
)	JURY TRIAL DEMANDED
COMMVAULT SYSTEMS, INC.,)	
)	
Defendant.)	

AMENDED COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff Rubrik, Inc. (“Rubrik”), by and through its attorneys, for its Amended Complaint against Defendant Commvault Systems, Inc. (“Commvault”), alleges as follows:

THE PARTIES

1. Rubrik is a Delaware corporation with a principal place of business at 3495 Deer Creek Road, Palo Alto, California 94304.
2. On information and belief, Commvault is a Delaware corporation with a principal place of business at 1 Commvault Way, Tinton Falls, New Jersey 07724.

JURISDICTION AND VENUE

3. This is a civil action for patent infringement arising under the patent laws of the United States, 35 U.S.C. §§ 1 *et seq.*
4. This Court has subject matter jurisdiction over the matters asserted herein under 28 U.S.C. §§ 1331 and 1338(a).
5. This Court has personal jurisdiction over Commvault in this action because, on information and belief, Commvault has committed and continues to commit acts of infringement

within the State of Texas and this judicial district giving rise to this action, including by, among other things, importing, offering to sell, and selling products and services that infringe the asserted patents.

6. Venue is proper in this judicial district under 28 U.S.C. § 1400(b). As noted above, on information and belief, Commvault has committed and continues to commit acts of infringement within this judicial district giving rise to this action, and Commvault has a regular and established place of business in this judicial district, including at 13620 Ranch Rd. 620, Suite A-150, Austin, Texas 78717.¹

FACTUAL BACKGROUND

Technology Overview

7. Modern businesses generate and retain vast amounts of data. The asserted patents relate to data backup and recovery, which is the process of creating and storing copies of data in such a way that data can be restored if needed, for example in the event of a data failure. Data failure can result from hardware or software failure, data corruption, human error, cyberattack, or some other event. Backup and recovery is an essential component of businesses' disaster recovery strategy, in addition to serving critical business needs with respect to data analytics.

Parties

8. Rubrik is an enterprise-level cloud data management company founded in January 2014. Since its inception, Rubrik has been repeatedly recognized as an innovative market disruptor in the \$48 billion cloud backup and recovery industry, and has enjoyed rapid success and adoption of its groundbreaking enterprise-grade data management solutions with consumer-grade simplicity.

¹ <https://www.commvault.com/contact-us>.

9. Rubrik has helped companies modernize and automate their data protection with its fast and scalable data management solutions. Rubrik's product portfolio includes its Cloud Data Management backup platform; Polaris, a SaaS-based platform; Mosaic for protection of NoSQL databases such as Apache Cassandra, DataStax Enterprise, and MongoDB, while assuring application availability; and Network-Attached Storage data management services DataProtect, DataDiscover, and DataFlow.

10. Commvault, a legacy player in the backup and recovery industry that was founded in the 1980s, also offers enterprise-level data management services.

Patents-in-Suit

11. U.S. Patent No. 10,852,998 (the "'998 patent"), entitled "Sub-Cluster Recovery Using a Partition Group Index," was duly and legally issued by the U.S. Patent and Trademark Office on December 1, 2020. The '998 patent is assigned to Rubrik. A copy of the '998 patent is attached as **Exhibit A**.

12. Rubrik owns all right, title, and interest in and to the '998 patent. Rubrik has the right to sue and recover for the infringement of the '998 patent.

13. U.S. Patent No. 9,075,773 (the "'773 patent"), entitled "Prioritized Repair of Data Storage Failures," was duly and legally issued by the U.S. Patent and Trademark Office on July 7, 2015. The '773 patent is assigned to Rubrik. A copy of the '773 patent is attached as **Exhibit B**.

14. Rubrik owns all right, title, and interest in and to the '773 patent. Rubrik has the right to sue and recover for the infringement of the '773 patent.

15. U.S. Patent No. 10,133,495 (the "'495 patent"), entitled "Converged Search and Archival System," was duly and legally issued by the U.S. Patent and Trademark Office on

November 20, 2018. The '495 patent is assigned to Rubrik. A copy of the '495 patent is attached as **Exhibit C**.

16. Rubrik owns all right, title, and interest in and to the '495 patent. Rubrik has the right to sue and recover for the infringement of the '495 patent.

17. The infringing Commvault products include, but are not limited to, the Commvault Complete Backup and Recovery, Commvault Data Protection, and Commvault Disaster Recovery software suites, related software features, and related products and services as identified and described in greater detail in Counts One, Two, and Three below, as well as any other Commvault products that practice the claims of the '998 patent, '773 patent, and/or the '495 patent (the "Accused Products").

COUNT ONE

(Infringement of United States Patent No. 10,852,998)


18. Rubrik incorporates herein and re-alleges the allegations of the preceding Paragraphs as if set forth fully herein.

19. The '998 patent claims methods and systems of enabling recovery of data in a partition group in a node of a distributed database through the use of a partition group index.

20. Apache Cassandra ("Cassandra") is an example of an open source, distributed database system that allows users to organize data within a cluster of multiple nodes.² These nodes are grouped together into a "cluster," in which data from one node is replicated to other nodes in the cluster:³

² <https://cassandra.apache.org/doc/latest/architecture/overview.html>

³ https://www.tutorialspoint.com/cassandra/cassandra_architecture.htm



21. Databases like Cassandra organize the data on a node into one or more containers or namespaces, known in Cassandra as KeySpaces. A Cassandra cluster contains one unique container (or KeySpace) per node.⁴ Data stored on one node is replicated to other nodes in the same cluster, and the containers define data replication between nodes.⁵ Therefore, in a typical Cassandra cluster, for example, each node will contain multiple KeySpaces due to replication.

22. In the context of distributed databases, such as Cassandra, the '998 patent claims systems and methods for recovery of individual partition groups (*e.g.*, data associated with particular KeySpaces) within a cluster by indexing backups of data tables to the partition group they originated from (*i.e.*, creating a partition group index).

23. The '998 patent discloses that, in a data storage environment having a plurality of nodes, scanning data items stored in the plurality of nodes to create a partition group index that indexes data items stored in the nodes into a plurality of partition groups, wherein each partition

⁴ https://www.tutorialspoint.com/cassandra/cassandra_create_keyspace.htm

⁵ <https://cassandra.apache.org/doc/latest/cql/ddl.html>

group corresponds to a node in the cluster and comprises a subset of data stored on the node. The '998 patent further discloses that nodes in a new or recovered cluster may be instantiated with data from particular partition groups by generating per node data for nodes in the new or recovered cluster based on mappings between the partition groups and the initial plurality of nodes represented in the partition group index. The '998 patent also discloses that data from one or more partition groups may be restored to a new node in the cluster from which it originated.

24. The '998 patent recites multiple technical implementation details that address the deficiencies in conventional prior art systems and methods for backup and restoration of distributed databases, namely the inability to restore data on a partition group basis. '998 patent at 1:14–16; 3:34–4:11. The implementation details and combination of steps in the '998 patent claims would be recognized by persons of skill in the art as unconventional.

25. For example, the '998 patent describes a novel and unconventional method of recovering subsets of data on a particular node in a cluster. For instance, prior art systems allowed for recovery on a full-node basis, but were not able to individually recover subsets of data (*i.e.*, partition groups, such as Cassandra KeySpaces) on a single node. *Id.* at 3:34–36; 3:60–4:1. The '998 patent improves upon these prior art systems and methods by retaining an index of data stored in a node on a partition group basis, where a partition group comprises a particular subset of data on the node, to enable retrieval and recovery of that specific partition group. *Id.* at 2:59–3:3; 3:6–4:1.

26. The '998 patent explains that “[u]sing the partition group index . . . data can be recovered on a partition group basis rather than on a physical database node basis.” *Id.* at 3:34–36. The unconventional creation and use of the partition group index as set forth in the '998 patent claims allows a data restoration system to “take[] the data items identified from the partition group

index for the partition group and fetch[] those data items from storage (*e.g.*, from a secondary data storage repository[,])” which can then be “formatted into a format (*e.g.*, recovery file) that can be understood by the database.” *Id.* at 3:43–51. By contrast, in prior art systems, “data items would have to be repopulated individually onto the database node by the database itself.” *Id.* at 3:63–4:1. The novel systems and methods claimed in the ’998 patent fix this problem by allowing data in any partition group—or combination of partition groups—to be “restored to a physical database node in one relatively quick file action.” *Id.* at 3:60–63. Recovery on a partition group basis has the additional benefit of improving overall system performance because “the production cluster will continue processing requests for other data item tables and will not get affected due to recovery of a particular data item table onto an independent node.” *Id.* at 4:8–11.

27. The ’998 patent further explains that the partition group index may be used to efficiently instantiate experimental clusters by selectively recovering data in only one or a few partition groups: “[f]or example, if an analytical team wants to run some experiment on production data backed up at a particular time[,], [t]he embodiments above allow the team to use an experimental cluster, not the production cluster, to experiment on that backed up data.” *Id.* at 4:17–21. In particular, “[u]sers can easily create a partition mapping (partition cluster to actual nodes) and, using the partition group index, efficiently create the database files” to selectively restore the data only in certain partition groups to a node in a new cluster. *Id.* at 4:25–30.

28. Commvault makes, uses, offers for sale, sells, and/or imports certain systems, products, and/or services in the United States, such as Commvault Complete Backup & Recovery (“Commvault Complete”), through which Commvault infringes, directly and indirectly, literally and/or under the doctrine of equivalents, at least claims 1, 8–10, 17, and 18 of the ’998 patent.

29. Commvault advertises that its Commvault Complete software is capable of restoring data associated with a specific KeySpace (*i.e.*, a partition group) within a Cassandra cluster. Upon information and belief, a Commvault Media Agent performs this operation by indexing backups (*e.g.*, snapshots) of data tables in a Cassandra cluster by the KeySpace that data is contained in (*i.e.*, the creation of a partition group index).

30. Commvault Complete is compatible with Cassandra. Commvault requires installation of a specialized Cassandra Agent on each node in a Cassandra cluster. The Cassandra Agent allows the cluster to interface with a Commvault Media Agent and thereby allow a Commvault user to “back up the entire Cassandra cluster environment using the Cassandra Agent.” In particular, the Cassandra Agent provides a “single interface to manage your Cassandra environment” that includes “full and incremental backup support” and allows “granular-level restores” of individual KeySpaces “to the source cluster or to a different Cassandra cluster from the source.”⁶

31. Commvault systems, products, and/or services satisfy all of the claim limitations of at least claims 1, 8–10, 17, and 18 of the ’998 patent. For example, independent claims 1 and 10 of the ’998 patent recite:

Claim 1. A method for sub-cluster recovery in a data storage environment having a plurality of nodes, the method comprising:

scanning data items stored in the first plurality of nodes of a first cluster;

while scanning, creating a partition group index, the partition group index indexing the data items into a plurality of partition groups, each partition group corresponding to a node of the first plurality of nodes and comprising a subset of data items stored in the node;

storing the index; and

instantiating a second cluster using a second plurality of nodes different, in number of nodes, from the first plurality of nodes, each node of the second plurality nodes being

⁶ <https://documentation.commvault.com/commvault/v11/article?p=30434.htm>

instantiated with data from at least one of the partition groups, the instantiating of the second cluster comprising for each node of the second plurality of nodes, generating per node data based on the mappings between the partition groups and the first plurality of nodes.

Claim 10. A system for sub-cluster recovery in a data storage environment having a plurality of nodes, the system comprising:

one or more non-transitory computer readable storage media;

a processing system operatively coupled with the one or more computer readable storage media; and

program instructions stored on the one or more computer readable storage media that, when read and executed by the processing system, direct the processing system to perform operations comprising:

scanning data items stored in the first plurality of nodes of a first cluster;

while scanning, creating a partition group index, the partition group index indexing the data items into a plurality of partition groups, each partition group corresponding to a node of the first plurality of nodes and comprising a subset of data items stored in the node;

storing the index; and

instantiating a second cluster using a second plurality of nodes different, in number of nodes, from the first plurality of nodes, each node of the second plurality nodes being instantiated with data from at least one of the partition groups, the instantiating of the second cluster comprising for each node of the second plurality of nodes, generating per node data based on the mappings between the partition groups and the first plurality of nodes.

32. A Commvault Cassandra Agent is installed and operates on each node in a Cassandra cluster, which constitutes a data storage environment having a plurality of nodes.⁷

33. A Commvault Media Agent, having read and write access to the Cassandra nodes, scans data items stored in database nodes for the purposes of deduplication. When a block of data is read from a source database, a unique signature is generated for that block using a hash algorithm. The Commvault Media Agent compares that signature against existing signatures for previously backed-up data blocks stored in a Deduplication Database. If the signature exists in the Deduplication Database, the block location is obtained from the associated Commvault Media

⁷ <https://documentation.commvault.com/commvault/v11/article?p=30436.htm>

Agent and this information is used in creating the object's index entry, while the duplicate data block is discarded. If the signature does not exist, the Media Agent writes the data block to destination storage and uses that location to create the data block's index entry in the Deduplication Database.⁸

34. On information and belief, the Commvault Media Agent creates a partition group index during the scanning of the nodes. During the deduplication process (*i.e.*, scanning), the Commvault Media Agent has write access to disk libraries where data blocks are stored, and tracks where blocks are written to destination storage in a Media Agent Index. This index allows the blocks to be reassembled so that the object can be restored or copied to other locations.⁹

35. In particular, on information and belief, the Commvault Media Agent creates a partition group index that indexes the data items scanned into a plurality of partition groups. Commvault advertises that its software “provides a granular browse option” when viewing backups of data in a Cassandra cluster, including the ability to “view the KeySpaces and the ColumnFamily entities [*i.e.* data] that KeySpace contains.”¹⁰ On information and belief, the Commvault Cassandra Agent associates scanned backup data (*e.g.*, a snapshot of a data table in a Cassandra node) with the KeySpace that data is associated with in an index, and stores this index, in order to allow a user to browse backup data by KeySpace and ColumnFamily.¹¹

⁸ <https://documentation.commvault.com/commvault/v11/article?p=12401.htm>

⁹ *Id.*

¹⁰ <https://documentation.commvault.com/commvault/v11/article?p=30475.htm>

¹¹ https://documentation.commvault.com/11.21/essential/107781_restore_cassandra_data.html

Restore Cassandra Data

Updated Monday, March 18, 2019

You can browse and restore backed up data for your subclient, individual keyspaces, or an entire instance.

Browsing and Restoring Cassandra Data Based on Backup Time

You can browse the Cassandra data either from the most recent backup or from a previous point in time.

- You can restore from the latest backups to keep your content on a secondary storage, such as tape, for long-term storage.
- You can restore from the data backed up on a specific date to restore to a point in time before the content became unusable.
- You can restore from the data backed for a date range if your content for that range was deleted accidentally.

36. A Commvault Media Agent can instantiate a second cluster of nodes different, in number of nodes, from the first plurality of nodes, where each node in the second cluster is instantiated with data from at least one of the partition groups. A Commvault Media Agent can perform an “out-of-place Cassandra restore” to “restore backed up data for individual KeySpaces,” wherein a “restore operation of the Cassandra KeySpace” can be conducted to new Cassandra cluster that has “[a] fewer number of nodes than the source” or “a larger number of nodes than the source.”¹² If the number of nodes in the source and destination clusters are not the same, the Commvault Cassandra Agent allows a user to set the same target node to multiple source nodes.¹³

37. A Commvault Media Agent instantiates each node in the new cluster with data by generating per node data based on the mapping between the partition groups and the first plurality of nodes. On information and belief, the Commvault Media Agent retrieves backups of data tables in the selected KeySpace(s) using the partition group index, which maps those data tables to that KeySpace. The Commvault Media Agent then generates per node data for the new cluster by

¹² <https://documentation.commvault.com/commvault/v11/article?p=30482.htm>; https://documentation.commvault.com/11.21/essential/107784_restoring_individual_keyspaces.html; *see also* <https://documentation.commvault.com/commvault/v11/article?p=30434.htm> (“Use granular-level restores to perform the following restore operations to the source cluster, or to a different Cassandra cluster from the source: . . . KeySpaces and ColumnFamily [or] a single node restore.”)

¹³ <https://documentation.commvault.com/commvault/v11/article?p=30482.htm>

configuring the retrieved backups into recovery files compatible with Cassandra. The Cassandra SSTableLoader tool is used to stream data across the nodes of the new cluster according to the replication strategy of the restored KeySpaces.¹⁴

38. Claims 8, 9, 17, and 18 of the '998 patent depend from either claim 1 or claim 10, and are similarly infringed by Commvault Complete when used to restore data associated with one or more individual KeySpaces in a Cassandra cluster. By way of example and without limitation, a Cassandra cluster comprises a NoSQL data store; and a Cassandra KeySpace comprises a group of data that is replicated across multiple nodes.

39. On information and belief, each of the above steps of the claimed methods is performed by Commvault in the United States, and each element of the claimed systems is found in Commvault's products and/or services that are made, used, offered for sale, sold, and/or imported into the United States. By making, using, offering for sale, selling, and/or importing into the United States Commvault's products and/or services without license or permission from Rubrik, Commvault has infringed, and will continue to infringe, the '998 patent in violation of 35 U.S.C. § 271(a).

40. In violation of 35 U.S.C. § 271(b), at least as of the filing of this Complaint, Commvault has also actively induced, and will continue to actively induce, users of its infringing system, products, and services to infringe the '998 patent. Commvault offered and continues to offer its infringing systems, services and products for sale, and instructed and continues to instruct users to operate them in an infringing manner through, without limitation, advertisements, product documentation or instructions, and customer support. By at least the filing of this Complaint, Rubrik has disclosed to Commvault the existence of the '998 patent and identified at least some

¹⁴ https://documentation.commvault.com/11.21/essential/107784_restoring_individual_keyspaces.html

of Commvault's systems, products, and/or services that infringe at least one claim of the '998 patent. Thus, based on this disclosure, Commvault has knowledge of the '998 patent and that its activities infringe the '998 patent. Based on Rubrik's disclosures, Commvault has also known or should have known since at least the filing of this Complaint that its customers, distributors, suppliers, and other purchasers of the Accused Products are infringing the '998 patent at least because Commvault has known that it is infringing the '998 patent. As a result of Commvault's inducement, users of its infringing system, products, and/or services have infringed and continue to infringe the '998 patent.

41. Commvault's infringement of the '998 patent has damaged and continues to damage Rubrik in an amount yet to be determined, of at least a reasonable royalty and/or the lost profits that Rubrik would have made but for Commvault's acts of infringement.

COUNT TWO

(Infringement of United States Patent No. 9,075,773)

42. Rubrik incorporates herein and re-alleges the allegations of the preceding Paragraphs as if set forth fully herein.

43. The '773 patent describes and claims methods and systems of prioritizing certain repair tasks over others based on characteristic values associated with those repair events.

44. The '773 patent discloses that data storage is managed over a network using a network computer that generates a repair task in response to a repair event associated with one or more new storage failures in the data storage system. Repair tasks are added to an ongoing task list, and prioritized based on a priority value determined for each repair task. The priority value is determined at least in part based on a comparison of one or more characteristic values associated with tasks in the list. The '773 patent also explains that a resource budget comprising network

bandwidth capacity is determined for a portion of the network, where the overall resource budget includes separate values for each of the different portions of the network. Repair tasks in the task list are promoted to active (*i.e.*, running) repair tasks based on their priority values and the resource budget: tasks with higher priority values are promoted over tasks with lower priority values, and as many tasks can run as active repair tasks as the resource budget permits.

45. The '773 patent recites technical implementation details that address the deficiencies in conventional prior art systems and methods for job management. The '773 patent explains that the amount of data and the size of storage systems used by modern enterprises has been increasing, and that “as the size of the data storage and the number of storage devices continue to increase, data restoration may increasingly take a disadvantageously long time to complete.” '773 patent at 1:23–26. Thus, the '773 patent improves upon prior art job management systems and methods by allowing resources such as network bandwidth capacity to be selectively allocated to the most time-critical repair tasks in a task list or job queue. The implementation details and combination of steps in the '773 patent claims would be recognized by persons of skill in the art as unconventional.

46. For example, the '773 patent provides an unconventional improvement to job management methods and systems that includes a novel method of determining a resource budget comprising network bandwidth capacity. The '773 patent explains that “bandwidth capacity for different portions and/or sub-networks may be determined and tracked independently as part of a resource budget.” '773 patent at 21:20–23. This unconventional method of determining a resource budget includes separate values for distinct portions of the network, such as “between the one or more storage computers, storage management server computers, intervening networks, or the like.” *Id.* at 21:16–19. The determined resource budget therefore allows “the topology of a network in

relationship to a particular storage failure [to] impact allocation of repair resources and/or the execution of active repair tasks.” *Id.* at 21:23–25.

47. The determination of the resource budget also allows a user “to modify and/or adjust one or more values in the resource budget and available for repair tasks and/or active repair tasks,” including “the network bandwidth that is made available for repair tasks.” *Id.* at 22:10–16. “For example, a user may limit the amount network bandwidth made available for repair resources to 50% of utilization to guarantee sufficient bandwidth for other users of the network.” *Id.* at 22:16–19.

48. The ’773 patent provides that this novel method of resource budget determination can be used to improve existing methods of prioritization of, and resource allocation between, competing repair tasks, particularly where different repair tasks have both different priorities and different network bandwidth requirements. Each repair task in a task list has a priority value associated with it which is determined by comparing characteristic values associated with different repair tasks.

49. The ’773 patent explains that repair tasks with higher priorities are promoted to active repair tasks and assigned resources from the resource budget to being executing. *Id.* at 3:35–49. In particular, “one or more repair tasks may be promoted to become active repair tasks based on the priority value of the repair tasks such that the promoted repair tasks have a higher priority . . . than other repair tasks in the task list, if any.” *Id.* at 4:41–44. Repair resources, such as network bandwidth capacity, “sufficient for execution are allocated to the one or more repair tasks that are promoted to active repair tasks.” *Id.* at 3:57–59. Furthermore, “if at least one other repair task may be available to use the unallocated repair resources, that repair task may be promoted to an active repair task [and] the unallocated repair resources may be allocated to the newly promoted

active repair task for execution.” *Id.* at 4:54–58. In this way, “[r]epair tasks may be prioritized and/or queued until they are allocated resources to perform actions, or otherwise promoted to be an active repair task.” *Id.* at 3:41–44.

50. Commvault makes, uses, offers for sale, sells, and/or imports certain systems, products, and/or services in the United States, such as the Job Manager and/or Job Controller functionalities of the Commvault CommCell Console (“Commvault Job Manager”), through which Commvault infringes, directly and indirectly, literally and/or under the doctrine of equivalents, at least claims 1, 7, 8, 14, and 15 of the ’773 patent.

51. Commvault systems, products, and/or services satisfy all of the claim limitations of at least claims 1, 7, 8, 14, and 15 of the ’773 patent. For example, independent claim 1 of the ’773 patent recites:

Claim 1. A method for managing data storage over a network using a network computer that executes instructions that perform actions, comprising:

if at least one repair event is associated with at least one new storage failure on a storage unit or a repair symbol unit, generating at least one new repair task that is associated with the at least one new storage failure, wherein the at least one new repair task is added to a task list that includes at least one other repair task for at least one current storage failure;

determining a priority value for each repair task in the task list based in part on a comparison of at least one characteristic value that is associated with each repair task in the task list, wherein a lower characteristic value generates a higher priority value over a higher characteristic value;

determining a resource budget based on a network bandwidth capacity for at least one different portion of the network, wherein the resource budget includes separate values for each of the different portions of the network;

promoting at least one repair task to be at least one new active repair task if the priority value for the at least one promoted repair task is higher than each other repair task and enough of the resource budget is available to execute the at least one new active repair task when each current working task is executing;

promoting at least one repair task to be the at least one new active repair task if a priority value for the at least one promoted task is higher than the priority level of at least one current active repair task, wherein each current active repair task is demoted to a repair task that has a lower priority value than the priority value for the at least one new active

repair task until enough of the resource budget is released to execute the at least one new active repair task;

promoting the at least one repair task to the at least one new active repair task if a remainder of the resource budget enables execution of the at least one new active repair task when the priority value of the at least one new active repair task is lower than at least one other active repair task requiring other than the remainder to execute; and

executing each active repair task to repair at least one storage failure that is associated with the active repair task.

52. The Commvault CommCell environment comprises one or more network computers that manage data storage over a network. “In a CommCell instance, the hardware components, such as the CommServe computer, client, MediaAgent, and storage devices (including libraries), are interconnected by a data network.”¹⁵ In particular, the Commvault Job Manager “manages and monitors immediate or scheduled operations throughout the CommCell environment.” The Commvault Job Manager communicates with Commvault Media Agents “to initiate and monitor data protection and data recovery operations,” “monitor resources necessary to complete these operations,” and “starts, suspends, and resumes operations based on these resources.”¹⁶

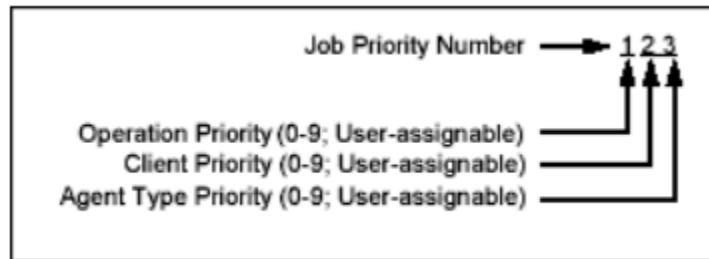
53. The Commvault Job Manager is capable of generating a repair task (*e.g.*, a restore task to recover a database from backup)¹⁷ and adding it to a task list that includes at least one other repair task. For example, the Commvault Job Manager maintains a “job queue” of active and pending repair tasks. Each repair task is assigned a three-digit job priority number by the

¹⁵ <https://documentation.commvault.com/commvault/v11/article?p=7068.htm>

¹⁶ <https://documentation.commvault.com/commvault/v11/article?p=3619.htm>

¹⁷ *E.g.*, <https://documentation.commvault.com/commvault/v11/article?p=12983.htm>

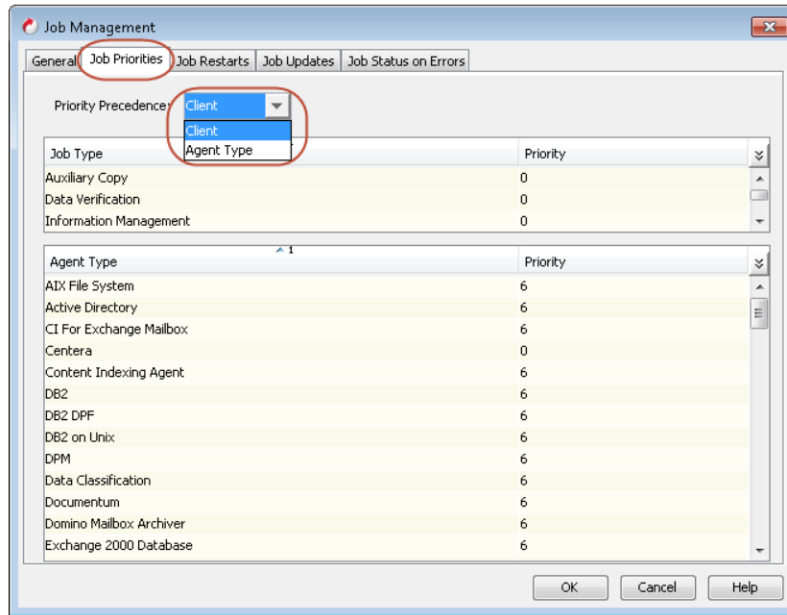
Commvault Job Manager. “The lower the job priority number is, the higher the priority assigned to the job, which determines its position in the job queue.”¹⁸



<https://documentation.commvault.com/commvault/v11/article?p=6540.htm>

54. The Commvault Job Manager determines a priority value for each repair task based on a comparison of characteristic values (*e.g.*, client priority and agent priority values) associated with each repair task. The first digit in a Commvault job priority value represents operation priority; for restores from backup, the default operation priority is 0 (*i.e.*, the highest-priority value). The second and third digits in a Commvault job priority value represent user-defined client and Media Agent priorities, and a user may determine which of these priority values is determinative. Therefore, for instance, if “client priority” is selected, the Commvault Job Scheduler will prioritize repair tasks with a lower client priority value (*i.e.*, “characteristic value”) over repair tasks with a higher client priority value.

¹⁸ <https://documentation.commvault.com/commvault/v11/article?p=6513.htm>



<https://documentation.commvault.com/commvault/v11/article?p=6516.htm>

55. The Commvault Job Manager determines a resource budget for repair tasks based on, for example, throttling thresholds set on network bandwidth capacity. “Network Bandwidth Throttling allows you to limit the amount of bandwidth that network-based operations, such as between clients and Media Agents, are allowed to consume.” Throttling may be employed to “set bandwidth to an absolute level or to a percentage of available bandwidth.” The determined resource budget includes separate values for each of the different portions of the network: throttling limits can be applied “for individual clients, client groups and Media Agents,” and “[o]nce configured, the throttling options are applied to all data transfer and control message operations, such as backup, copy, and restore.”¹⁹

56. The Commvault Job Manager is capable of promoting a repair task to an active repair task when its priority value is higher than each other non-active repair task when there are enough available resources to execute the new active repair tasks without demoting a current active

¹⁹ <https://documentation.commvault.com/commvault/v11/article?p=7483.htm>

repair task. The Commvault Job Manager prioritizes promotion of repair tasks such that, “if there is contention for resources, the job with the highest priority gets the resources first.”²⁰

57. The Commvault Job Manager is also capable of promoting a repair task to an active repair task when its priority value is higher than at least one other active repair task, and demoting the other repair tasks with lower priority values until enough resources are made available to execute the new, higher-priority active repair task. In particular, the Commvault Job Manager “job preemption,” where a higher priority job can preempt a lower priority job that is currently active.²¹ “When several jobs have different priorities and if Job Preemption is allowed, then the Job Manager interrupts the running job and allocates the freed resources to a higher priority job.”²² “The interrupted job enters a waiting state and resumes when the resources it needs become available.”²³

58. The Commvault Job Manager is capable of promoting a repair task to an active repair task if the new active repair task has a lower priority than another active repair task if there are adequate resources to execute the new active repair task. If multiple existing repair tasks are not competing for resources, “sources are allocated on a first-come, first-served basis.” The Commvault Job Manager then executes each active repair task.

59. On information and belief, each of the above steps of the claimed method is performed by Commvault in the United States, and each element of the claimed system is found in Commvault’s products and/or services that are made, used, offered for sale, sold, and/or

²⁰ <https://documentation.commvault.com/commvault/v11/article?p=6513.htm>

²¹ <https://documentation.commvault.com/commvault/v11/article?p=6540.htm>

²² <https://documentation.commvault.com/commvault/v11/article?p=6540.htm>

²³ <https://documentation.commvault.com/commvault/v11/article?p=6346.htm>

imported into the United States. By making, using, offering for sale, selling, and/or importing into the United States Commvault's products and/or services without license or permission from Rubrik, Commvault has infringed, and will continue to infringe, the '773 patent in violation of 35 U.S.C. § 271(a).

60. In violation of 35 U.S.C. § 271(b), at least as of the filing of this Complaint, Commvault has also actively induced, and will continue to actively induce, users of its infringing system, products, and services to infringe the '773 patent. Commvault offered and continues to offer its infringing systems, services and products for sale, and instructed and continues to instruct users to operate them in an infringing manner through, without limitation, advertisements, product documentation or instructions, and customer support. By at least the filing of this Complaint, Rubrik has disclosed to Commvault the existence of the '773 patent and identified at least some of Commvault's systems, products, and/or services that infringe the '773 patent. Thus, based on this disclosure, Commvault has knowledge of the '773 patent and that its activities infringe the '773 patent. Based on Rubrik's disclosures, Commvault has also known or should have known since at least the filing of this Complaint that its customers, distributors, suppliers, and other purchasers of the Accused Products are infringing the '773 patent at least because Commvault has known that it is infringing the '773 patent. As a result of Commvault's inducement, users of its infringing system, products, and/or services have infringed and continue to infringe the '773 patent.

61. Commvault's infringement of the '773 patent has damaged and continues to damage Rubrik in an amount yet to be determined, of at least a reasonable royalty and/or the lost profits that Rubrik would have made but for Commvault's acts of infringement.

COUNT THREE

(Infringement of United States Patent No. 10,133,495)

62. Rubrik incorporates herein and re-alleges the allegations of the preceding Paragraphs as if set forth fully herein.

63. U.S. Patent No. 10,133,495 (the “’495 patent”) claims systems and methods of identifying and restoring a particular version of a file from a portion of a version of a virtual machine using a virtual machine search index.

64. As explained in the ’495 patent, “[v]irtualization allows virtual hardware to be created and decoupled from the underlying physical hardware. For example, a hypervisor running on a host machine or server may be used to create one or more virtual machines.” ’495 patent, 1:13–16. These virtual machines “comprise a software implementation of a physical machine,” complete with “virtual memory, a virtual disk,” and the ability to run “a Windows® operating system” or “a Unix-like operating system such as OS X®.” *Id.* at 1:18–25. In this way, virtual machines mimic many of the key features of conventional, physical computers.

65. One common method of preserving the state of (*i.e.*, backing up) a virtual machine is through a snapshot backup. A snapshot backup is a type of backup copy that captures the state of an application, disk or system a particular point in time. The snapshot may then be used to restore a virtual machine to the state captured by the snapshot.²⁴

66. The ’495 patent discloses systems and methods of restoration of a particular version of a file associated with a portion of a version of a virtual machine. In particular, the ’495 patent discloses that “a particular version of a file to be restored is identified,” which “may

²⁴ <https://docs.vmware.com/en/VMware-Fusion/12/com.vmware.fusion.using.doc/GUID-E4C0FD3C-B527-42D3-A60A-B4249A43CD5F.html>.

comprise or correspond with a database, a spreadsheet, a word processing document, an image file, a video file, a text file, an executable file, an audio file, an electronic message, or an email.” ’495 patent, 32:54–59. The file “may be stored on a virtual disk of a virtual machine.” *Id.* at 32:55–56. “The particular version of the file may be selected by an end user of a storage appliance . . . using a user interface provided by the storage appliance.” *Id.* at 32:59–62. Then, the ’495 patent explains that a virtual machine search index is “acquired” and used to identify “a version of the virtual machine that includes the particular version of the file.” *Id.* at 32:62–67. “[A] base image for generating the version of the virtual machine is identified,” along with a “set of incremental files for generating the version of the virtual machine.” *Id.* at 33:1–6. The ’495 patent discloses that “a first portion of the base image that includes file system metadata for the virtual disk storing the file is determined,” and “a portion of the version of the virtual machine is generated using the file system metadata and the set of incremental files.” *Id.* at 33:12–20. Finally, “the particular version of the file is extracted from the portion of the version of the virtual machine,” and “the particular version of the file is outputted.” *Id.* at 33:20–24.

67. The ’495 patent recites multiple technical implementation details that address the deficiencies in conventional prior art systems and methods for point-in-time recovery of particular versions of files from backups of virtual machines. In particular, the ’495 patent describes novel and unconventional systems and methods of identification and restoration of a particular file within a version of a virtual machine through the use of a “virtual machine search index” that allows searching a versions of virtualized machines based on metadata associated with the particular file. *See, e.g.*, ’495 patent, 14:21–15:3, 31:31–33:58, Figs. 5A–5D. The ’495 patent explains that “the virtual machine search index 106 may be used to identify a version of the virtual machine that

includes a particular version of the file (*e.g.*, a particular version of a database, a spreadsheet, or a word processing document).” *Id.* at 14:29–32.

68. The ’495 patent additionally improves on prior art systems and methods by allowing for restoration of only a portion of the identified version of the virtual machine. The ’495 patent explains that “[o]ne benefit of extracting the particular version of the file from a portion of the version of the virtual machine that includes the particular version of the file is that the particular version of the file may be quickly restored without having to first restore an entire image of the version of the virtual machine.” *Id.* at 33:30–35. The implementation details and combination of steps in the ’495 patent claims would be recognized by persons of skill in the art as unconventional.

69. Commvault makes, uses, offers for sale, sells, and/or imports certain systems, products, and/or services in the United States, such as Commvault Complete Backup & Recovery, through which Commvault infringes, directly and indirectly, literally and/or under the doctrine of equivalents, at least claims 1–4 and 12–15 of the ’495 patent.

70. In particular, Commvault Complete Backup & Recovery software, when installed and operating in a data management system (the “Commvault Backup System”), is advertised as allowing a user to “view[] and restor[e] a file version,” including the ability to “view multiple versions of a file that has changed across different backup jobs” of a virtual machine.²⁵ On information and belief, the Commvault Backup System, including but not limited to the Commvault VMware Agent, is capable of restoring an identified version of a file by restoring only a portion of an identified version of a virtual machine.

²⁵ <https://documentation.commvault.com/commvault/v11/article?p=36510.htm>

71. Commvault systems, products, and/or services satisfy all of the claim limitations of at least 1–4 and 12–15 of the '495 patent. For example, independent claims 1 and 12 of the '495 patent recite:

Claim 1. A method for operating a data management system, comprising:

identifying a particular version of a file to be restored;

acquiring a virtual machine search index for a virtual machine, the virtual machine search index includes a mapping of the particular version of the file to a point in time snapshot of the virtual machine that includes the particular version of the file;

identifying a version of the virtual machine that includes the particular version of the file using the virtual machine search index;

identifying a base image for generating a portion of the version of the virtual machine;

identifying a set of incremental files for generating the portion of the version of the virtual machine;

determining a first portion of the base image that includes file system metadata, the file system metadata identifies a location of the file to be restored within the version of the virtual machine;

generating the portion of the version of the virtual machine less than the entire version of the virtual machine using the file system metadata and the set of incremental files, the generating the portion of the version of the virtual machine includes determining one or more regions within the base image that store the file to be restored and patching the one or more regions using the set of incremental files;

extracting the particular version of the file from the portion of the version of the virtual machine, the extracting the particular version of the file includes extracting the particular version of the file from the portion of the version of the virtual machine without generating a full image of the version of the virtual machine; and

outputting the particular version of the file.

Claim 12. A data management system, comprising:

a memory configured to store a base image for generating a version of a virtual machine, a set of incremental files for generating the version of the virtual machine, and a virtual machine search index for the virtual machine; and

one or more processors in communication with the memory, the virtual machine search index includes a mapping of a particular version of a data object to a point in time snapshot of the virtual machine that includes the particular version of the data object, the one or

more processors configured to identify the particular version of the data object to be restored and determine that the version of the virtual machine includes the particular version of the data object using the virtual machine search index, the one or more processors configured to determine a first portion of the base image that includes metadata for locating the data object within a virtual disk of the virtual machine, the one or more processors configured to generate a portion of the version of the virtual machine less than the entire version of the virtual machine using the metadata and the set of incremental files, the one or more processors configured to determine one or more regions within the base image that store the data object to be restored and patch the one or more regions using the set of incremental files, the one or more processors configured to extract the particular version of the data object from the portion of the version of the virtual machine without generating a full image of the version of the virtual machine and output the particular version of the data object.

72. The Commvault Backup System identifies a particular version of a file to be restored. Commvault advertises a “Restores for VMware Intellisnap” feature that allows users to restore “full VMs, VMDK files, or guest files and folders.”²⁶ In particular, the Commvault Backup System has a “View All Versions” feature that “enable [a user] to view multiple versions of a file that has changed across different backup jobs.” A Commvault user may “[u]se this feature to restore a specific version of a file from a Browse and Restore or Find Operation.”²⁷

73. The Commvault Backup System acquires a virtual machine search index for a virtual machine. “Using the Command Center, [a Commvault user] can file index virtual machines (VMs).” Commvault advertises that “[f]ile indexing is the process of searching full, incremental, or synthetic full backups in order to create a file (called an index), which contains all of the metadata (such as file name and date) that Commvault needs to be able to locate and perform operations on the data that it backs up (for example, restore data).” “This file indexing (for

²⁶ <https://documentation.commvault.com/commvault/v11/article?p=36454.htm>

²⁷ <https://documentation.commvault.com/commvault/v11/article?p=36510.htm>

example, on VM-level data) results in faster file-level browse operations.” By default, the Commvault Backup System performs file indexing after a backup is completed.²⁸

74. On information and belief, the virtual machine search index acquired by the Commvault Backup System includes a mapping of the particular version of the file to a point in time snapshot of the virtual machine that includes the particular version of the file. For example, the Commvault Backup System generates an file index by performing a “file indexing job,” which includes launching a “Filescan operation” on “each VM backup.” “The Filescan operation processes each volume to create a list of files and folders with the following attributes: Disks and volumes, File and folder names, File size, [and] File type.”²⁹

75. The Commvault Backup System identifies a version of the virtual machine that includes the particular version of the file using the virtual machine search index. Commvault advertises that a Commvault user may “execute a Find operation from a backup set or subclient,” by “enter[ing] the file name or pattern in the File Name box” and “in the Search within these folders box, enter[ing] the location where [the user] want[s] to search for the file.” Then, “the file that match the criteria are displayed in the Search Results tab”; the user may “select a file on Search Results tab[,] then right-click the file and select View All Versions.” “The available versions for the selected file are displayed in the right pane,” and the user may select the file version to restore.³⁰

76. The Commvault Backup System identifies a base image for generating a portion of the version of the virtual machine. Commvault’s snapshot feature, IntelliSnap, supports different types of snapshots, including full backups and incremental backups. “A full backup contains all

²⁸ <https://documentation.commvault.com/commvault/v11/article?p=101637.htm>

²⁹ <https://documentation.commvault.com/commvault/v11/article?p=101917.htm>

³⁰ <https://documentation.commvault.com/commvault/v11/article?p=36510.htm>

the data in the subclient contents,” and is therefore “a baseline to which subsequent backup types are applied.”³¹

77. On information and belief, the Commvault Backup System identifies a set of incremental files for generating the portion of the version of the virtual machine. Commvault’s IntelliSnap for VMware supports incremental backups. “During incremental backups, Changed Block Tracking is used to optimize backups of virtual machines by reading only the allocated and modified portions of a virtual disk.”³²

78. The Commvault Backup System determines a first portion of the base image that includes file system metadata, wherein the file system metadata identifies a location of the file to be restored within the version of the virtual machine. Commvault advertises that file indexing for virtual machines “is the process of searching full, incremental, or synthetic full backups in order to create a file (called an index), which contains all of the metadata (such as file name and date) that Commvault needs to be able to locate and perform operations on the data that it backs up (for example, restore data).” A “Key Feature” of this functionality is a user’s ability to “automatically index metadata for files and folders,” and then use that index to “search virtual machine files and folders by file name or file extension.”³³

79. On information and belief, the Commvault Backup System generates the portion of the version of the virtual machine less than the entire version of the virtual machine using the file system metadata and the set of incremental files, and determines one or more regions within the base image that store the file to be restored and patching the one or more regions using the set of

³¹ <https://documentation.commvault.com/commvault/v11/article?p=11684.htm>.

³² <https://documentation.commvault.com/commvault/v11/article?p=36407.htm>

³³ <https://documentation.commvault.com/commvault/v11/article?p=101637.htm>

incremental files. Commvault advertises that a Commvault user may “restore guest files and folders from a virtual machine streaming backup or IntelliSnap backup to a virtual machine or destination client.” To do so, the Commvault user may “right-click the subclient that contains the data to restore and select Browse and Restore,” “select Guest Files and Folders,” and “navigate to the required volume on the virtual machine with the files and folders to be restored,” after which “the guest files and folders on the volume are displayed in the right pane.” The Commvault user may then “select the required guest files and folders in the right pane,” including “selecting a specific file version” from the “View All Versions” tab, and “click Recover All Selected.”³⁴ Commvault alternatively instructs users that, “to recover larger files or in other cases where you cannot restore a file, restore the complete virtual machine or the disk that contains the file.”³⁵

80. The Commvault Backup System extracts a particular version of the file from the portion of the version of the virtual machine, the extracting the particular version of the file includes extracting the particular version of the file from the portion of the version of the virtual machine without generating a full image of the version of the virtual machine. In addition to the foregoing, Commvault advertises that a Commvault user may specify “restore options” including to the proxy (*i.e.* location) to which the files are restored, or for agentless restores, the specific virtual machine to which the files are restored.³⁶

81. Finally, the Commvault Backup System outputs the particular version of the file. Commvault advertises that “[a] Commvault user can restore guest files and folders from a virtual machine streaming backup or IntelliSnap backup to a virtual machine or destination client.”

³⁴ <https://documentation.commvault.com/commvault/v11/article?p=36510.htm>

³⁵ <https://documentation.commvault.com/commvault/v11/article?p=36492.htm>

³⁶ <https://documentation.commvault.com/commvault/v11/article?p=36509.htm>

82. Dependent claims 2–4 and 13–15 are similarly infringed by the Commvault Backup System. By way of example and without limitation, the Commvault Backup System is capable of identifying a plurality of regions within a base image to restore, and can transmit the particular version of the identified file to a computing device. Furthermore, the Commvault Backup System is capable of acquiring a merged file corresponding with the version of the virtual machine and identifying the base image and set of incremental files using the merged file.

83. On information and belief, each of the above steps of the claimed methods is performed by Commvault in the United States, and each element of the claimed systems is found in Commvault’s products and/or services that are made, used, offered for sale, sold, and/or imported into the United States. By making, using, offering for sale, selling, and/or importing into the United States Commvault’s products and/or services without license or permission from Rubrik, Commvault has infringed, and will continue to infringe, the ’495 patent in violation of 35 U.S.C. § 271(a).

84. In violation of 35 U.S.C. § 271(b), at least as of the filing of this Complaint, Commvault has also actively induced, and will continue to actively induce, users of its infringing system, products, and services to infringe the ’495 patent. Commvault offered and continues to offer its infringing systems, services and products for sale, and instructed and continues to instruct users to operate them in an infringing manner through, without limitation, advertisements, product documentation or instructions, and customer support. By at least the filing of this Complaint, Rubrik has disclosed to Commvault the existence of the ’495 patent and identified at least some of Commvault’s systems, products, and/or services that infringe at least one claim of the ’495 patent. Thus, based on this disclosure, Commvault has knowledge of the ’495 patent and that its activities infringe the ’495 patent. Based on Rubrik’s disclosures, Commvault has also known or

should have known since at least the filing of this Complaint that its customers, distributors, suppliers, and other purchasers of the accused products are infringing the '495 patent at least because Commvault has known that it is infringing the '495 patent. As a result of Commvault's inducement, users of its infringing system, products, and/or services have infringed and continue to infringe the '495 patent.

85. Commvault's infringement of the '495 patent has damaged and continues to damage Rubrik in an amount, yet to be determined, of at least a reasonable royalty and/or the lost profits that Rubrik would have made but for Commvault's acts of infringement.

JURY DEMAND

86. Pursuant to Rule 38(b) of the Federal Rules of Civil Procedure, Rubrik respectfully requests a trial by jury on all issues so triable.

REQUESTED RELIEF

WHEREFORE, Rubrik respectfully seeks the following relief:

- a) The entry of judgment declaring that Commvault has infringed the '998, '773, and '495 patents;
- b) An order that, in accordance with 35 U.S.C. § 283, Commvault and all its affiliates, employees, agents, officers, directors, attorneys, successors, and assigns and all those acting on behalf of or in active concert or participation with any of them, be permanently enjoined from (1) infringing the '998, '773, '495 patents and (2) making, using, selling, offering for sale and/or importing the accused products described in Counts One, Two, and Three;
- c) An award of all available damages for Commvault's infringement of the '998, '773, and '495 patents, not less than a reasonable royalty, together with pre-judgment and post-judgment interest;

d) An order requiring Commvault to provide a pre-judgment accounting and to pay supplemental damages to Rubrik, including without limitation, pre-judgment and post-judgment interest;

e) The entry of an order declaring that this is an exceptional case and awarding Rubrik its costs, expenses, and reasonable attorney fees under 35 U.S.C. § 285 and all other applicable statutes, rules, and common law; AND

f) An order awarding Rubrik any such further relief as the Court may deem just and proper under the circumstances.

DATE: May 6, 2021

Respectfully submitted,

By: /s/ Josh Krevitt

Josh Krevitt
Brian A. Rosenthal
Paul E. Torchia
Katherine Q. Dominguez
Florina Yezril
Allen Kathir
Kyanna Sabanoglu
GIBSON, DUNN & CRUTCHER LLP
200 Park Avenue
New York, NY 10166
Tel: (212) 351-4000

Jaysen S. Chung
Ryan K. Iwahashi
Christine L. Ranney
GIBSON, DUNN & CRUTCHER LLP
555 Mission Street, Suite 3000
San Francisco, CA 94015
Tel: (415) 393-8200

Barry K. Shelton
SHELTON COBURN LLP
bshelton@sheltoncoburn.com
Texas State Bar No. 24055029
311 RR 620, Suite 205
Austin, TX 78734-4775
Tel: (512) 263-2165

*Attorneys for Plaintiff and
Counterclaim-Defendant Rubrik, Inc.*

CERTIFICATE OF SERVICE

The undersigned hereby certifies that a true and correct copy of the above and foregoing document has been served on all counsel of record via the Court's ECF system.

/s/ Barry K. Shelton

Barry K. Shelton